<EURO/SYS'20>

AlloX: Compute Allocation in Hybrid Clusters

Tan N. Le Xiao Sun Mosharaf Chowdhury Zhenhua Liu

Stony Brook University

SUNY Korea The State University of New York

UNIVERSITY OF MICHIGAN

tnle@cs.stonybrook.edu

Resource Allocation in Clusters

Performance

Fairness

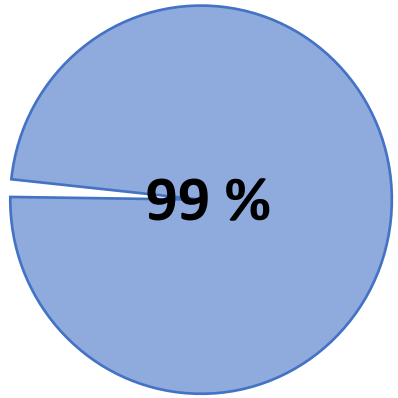

Utilization

Resource Allocation Design Space

Single Resource

Time

CPU Sharing

Space

Memory Sharing

Multiple Resources

Traditional

DRF (nsdi'11),
Carbyne (osdi'16)
HUG (nsdi'16)

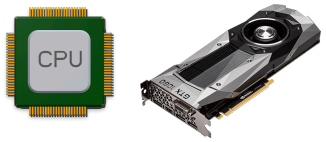

Interchangeable

Deadlines

TetriSched
(eurosys'16)

Performance & Fairness

AlloX

Interchangeability in Resources

Same applications run on different resource types

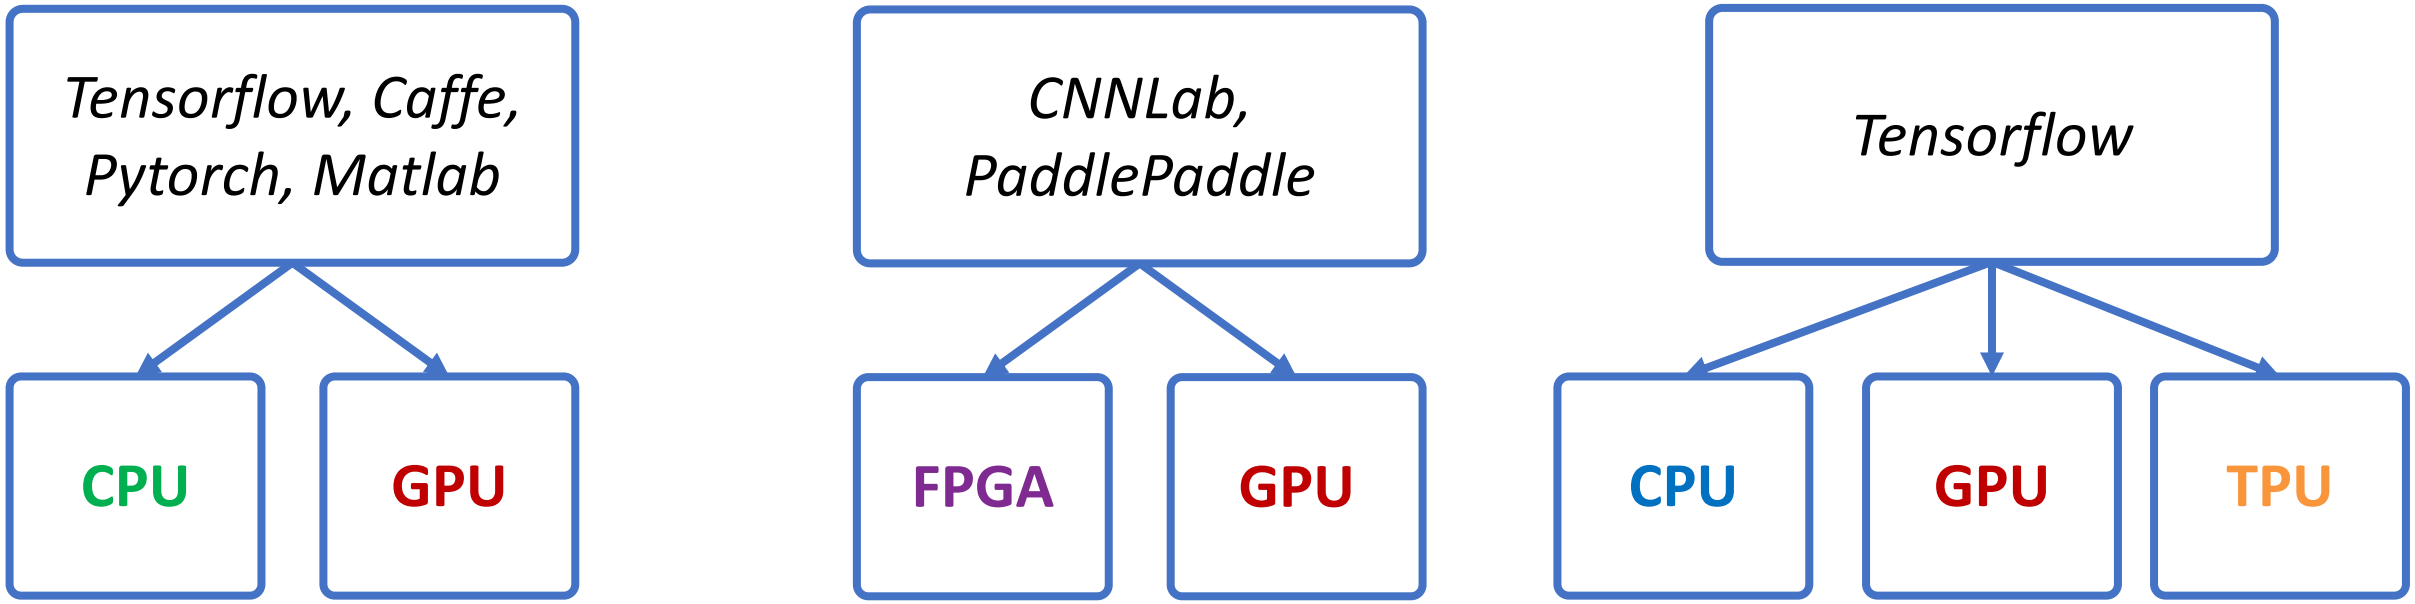

Modern Frameworks support Interchangeability

https://github.com/PaddlePaddle/Paddle
https://github.com/cnnlabs

Heterogeneity in hybrid CPU/GPU Clusters

Traditional nodes

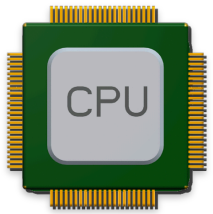

Expensive GPUs

Speed-up rates are distinct

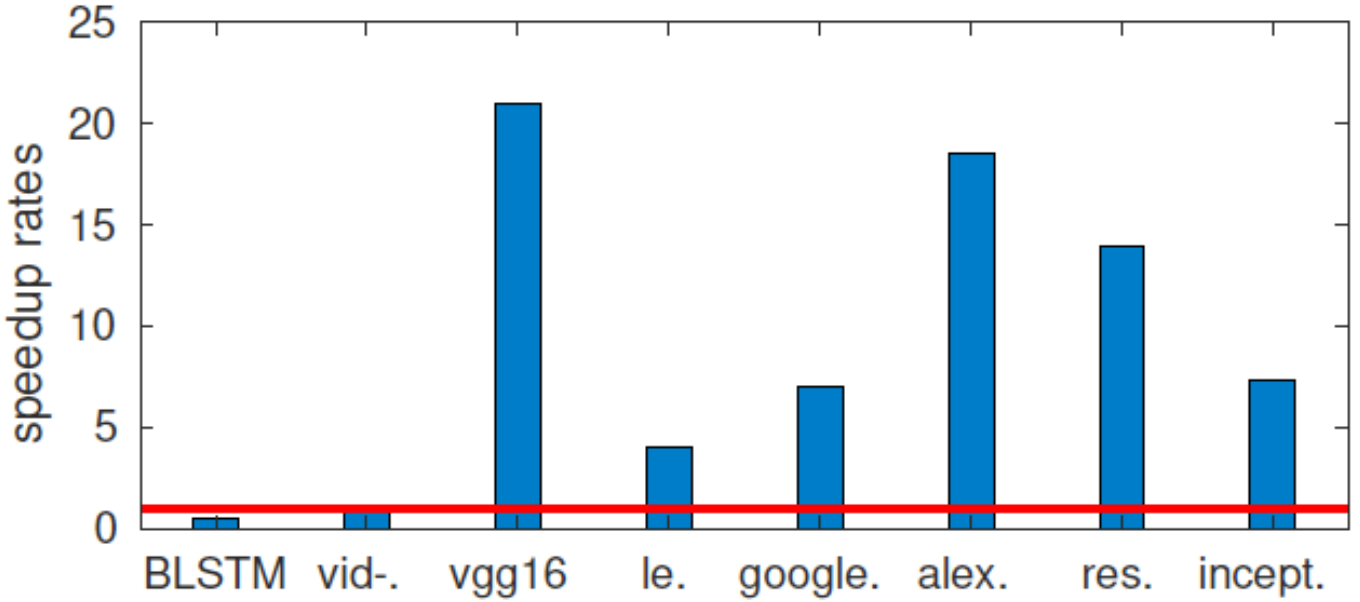

Intel E5 2.4Ghz CPU vs. Nvidia K80 GPU

Overload if most users prefer GPUs

Expensive GPUs are overloaded while **CPUs are under-utilized**

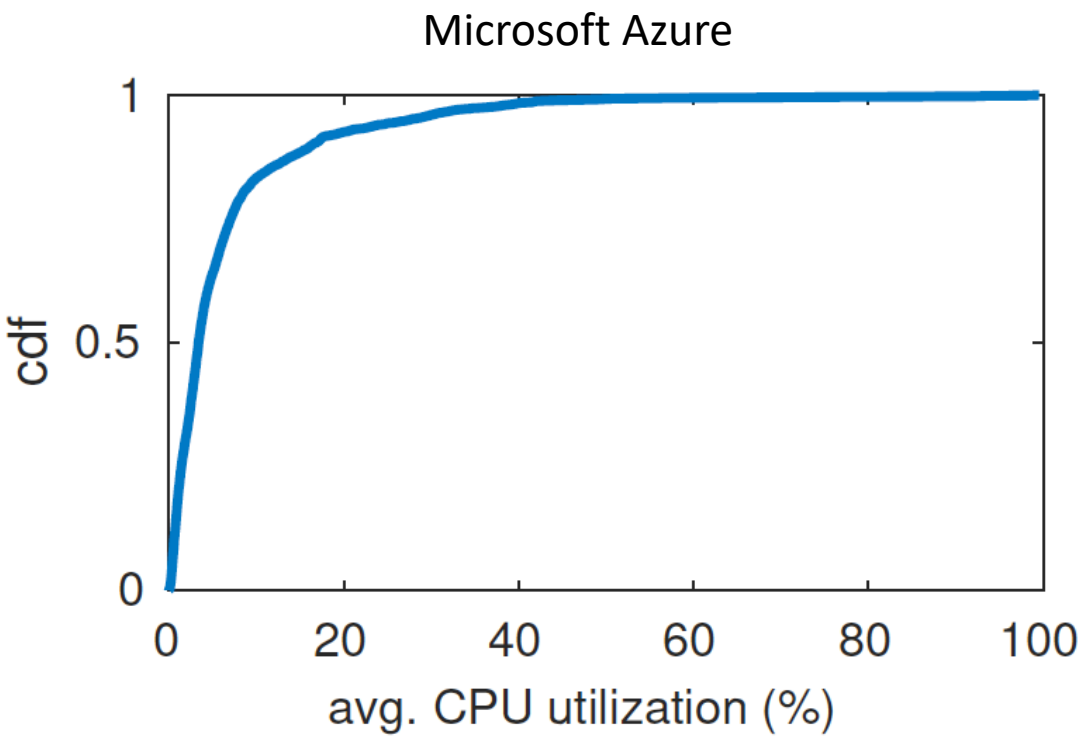

Let's explore some solutions

Join the Shortest Queue (JSQ)

Processing times
(GPU, CPU)

J1 (**40**, 50)
J2 (**30**, 40)
J3 (35, **150**)
J4 (50, **160**)

JSQ

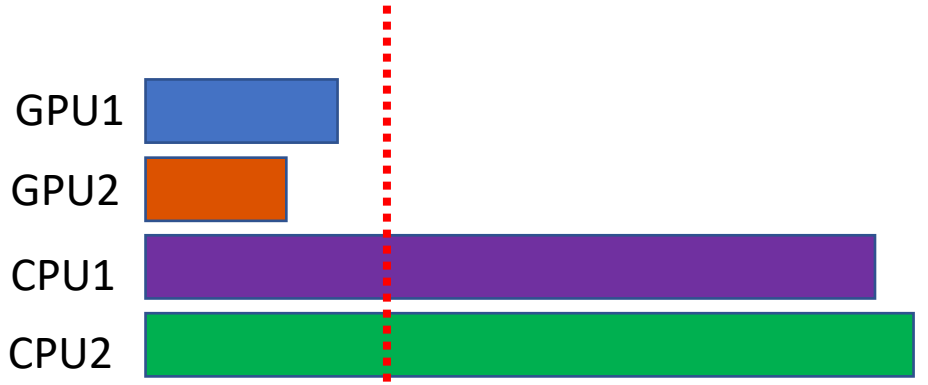

-69% makespan

Optimal

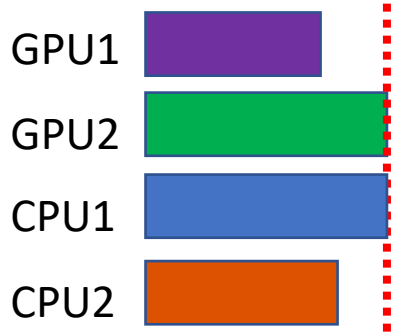

-54% avg. compl. time

JSQ does not consider processing times

Shortest Job First (SJF)

Processing times
(GPU, CPU)

J1 (**10**, 20)
J2 (**15**, 25)
J3 (20, **100**)
J4 (20, **90**)

SJF

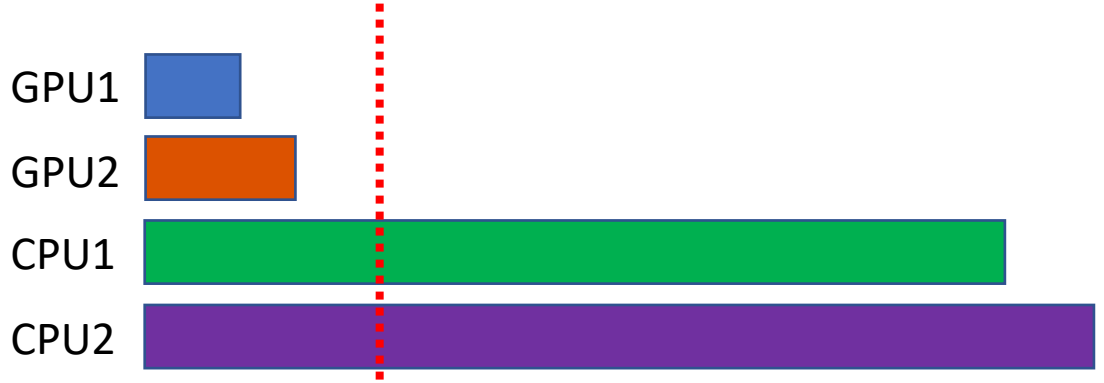

Optimal

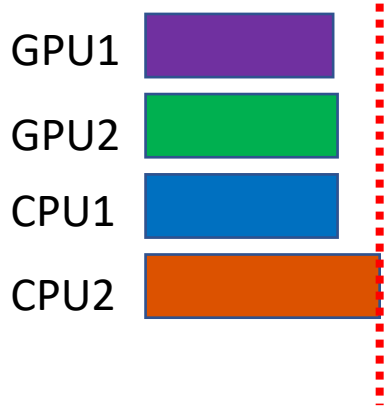

-75% makespan

-60% avg. compl. time

SJF does not consider speed-up rates

AlloX – Minimize Avg. Completion Time

Convert the scheduling & placement

into min-cost bipartite matching

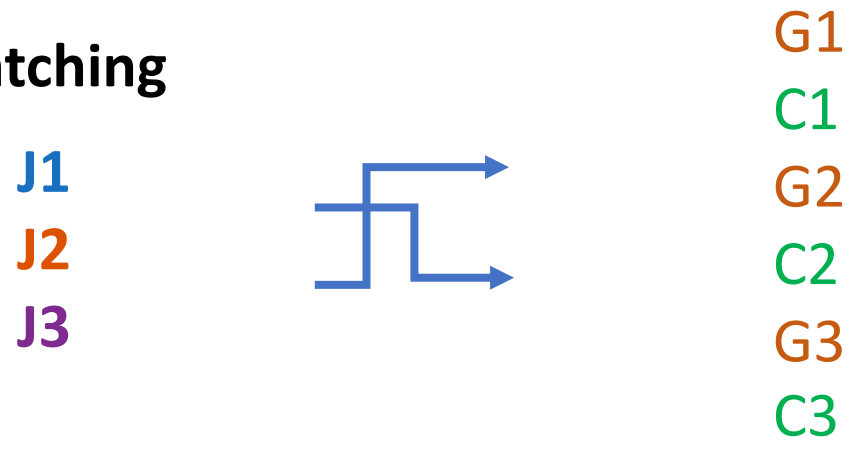

solved in polynomial time

AlloX – Maintains Fairness for interchangeable resources

User A may not be happy if we keep putting him on CPU.

Idea: Prioritize users with low fairness scores *F*
who run jobs on the unfavorable resources

User 1 F-1
User 2 F-2

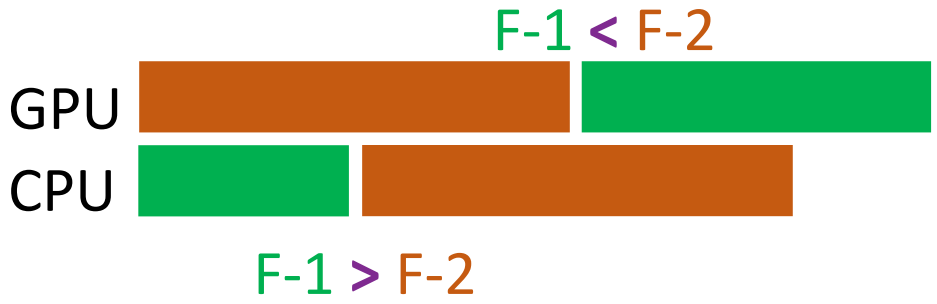

AlloX System

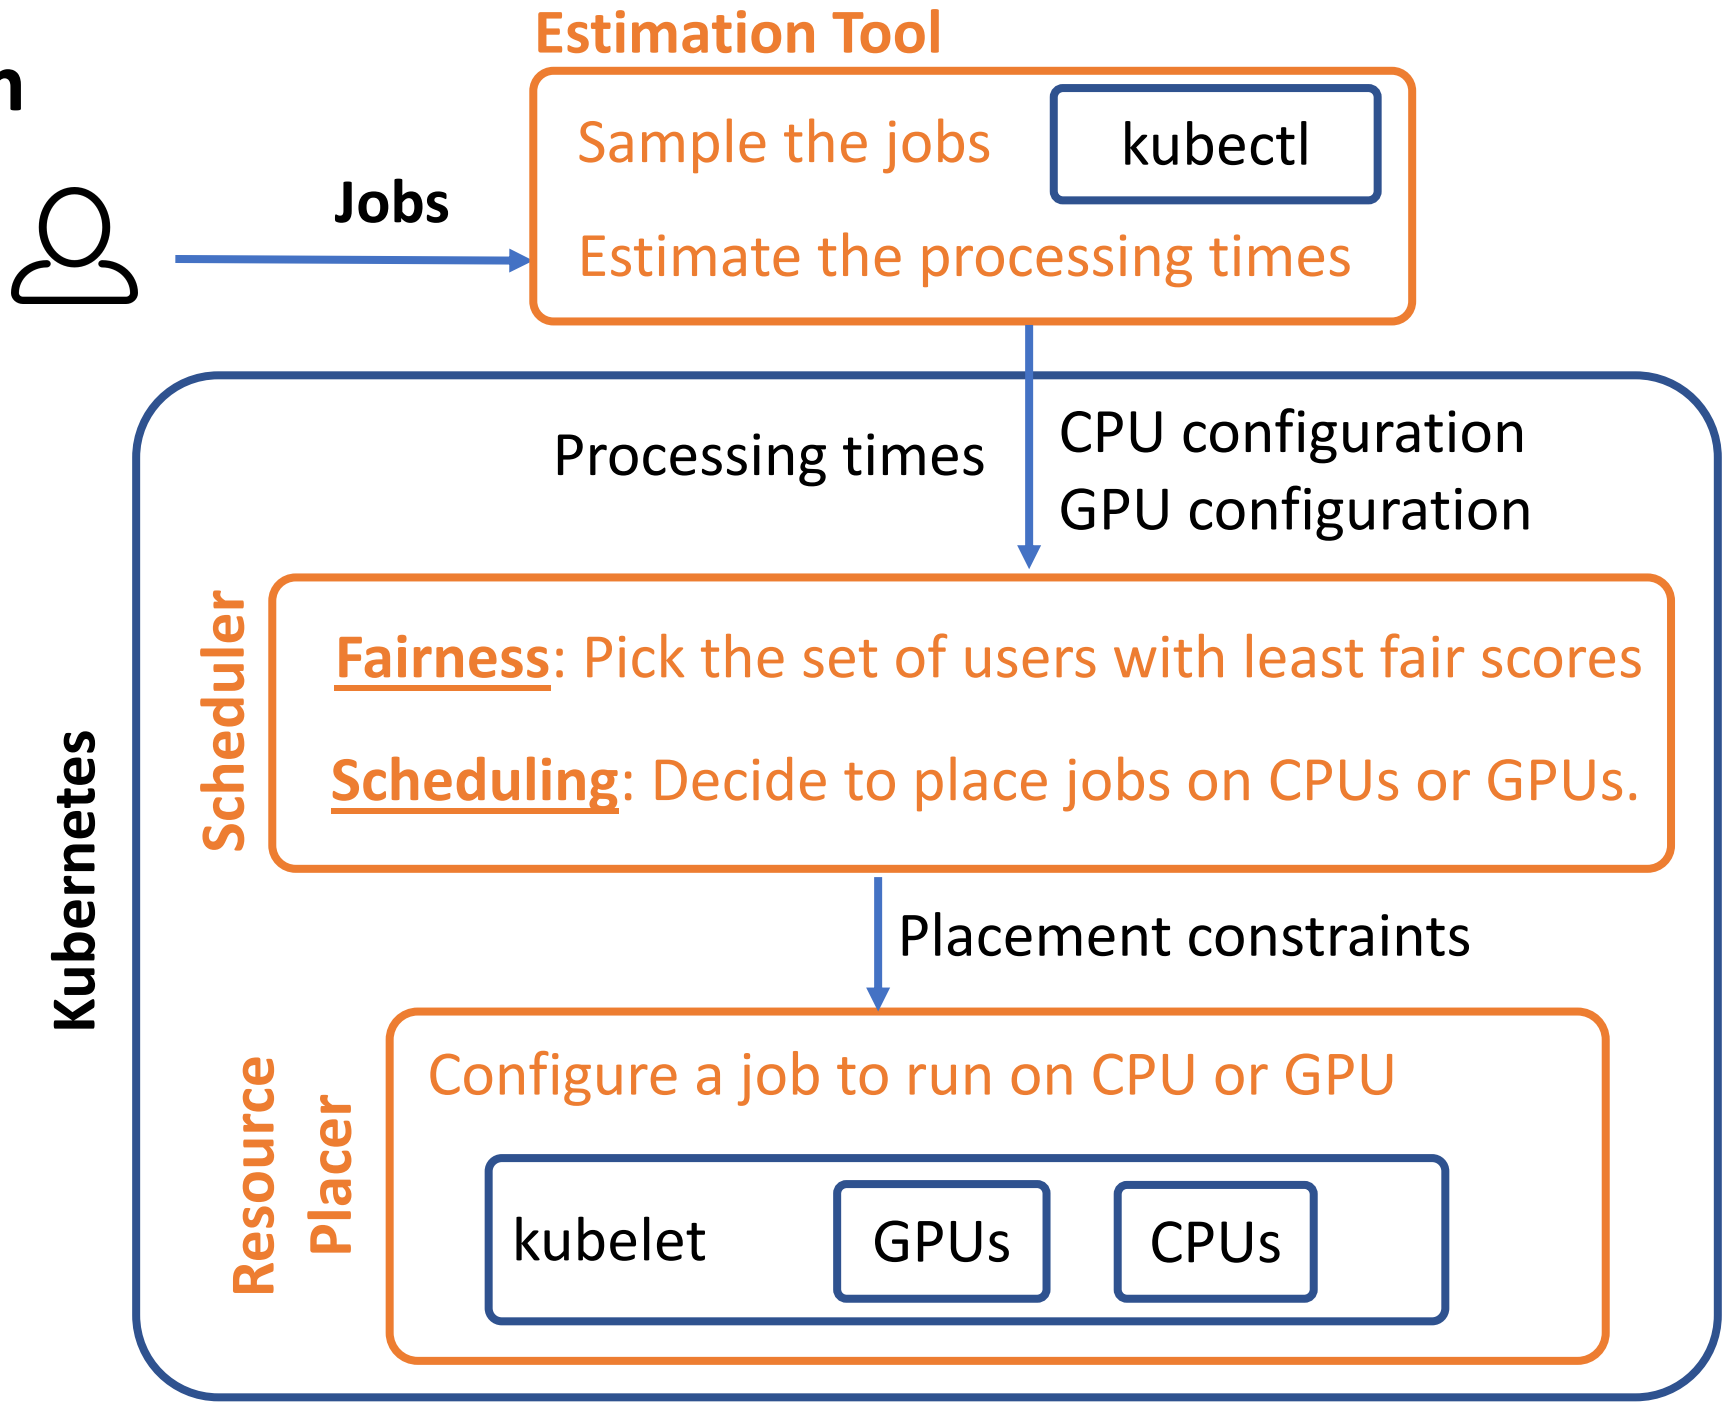

Performance of AlloX

DRF: Dominant Resource Fairness + FIFO
Resource configurations are fixed

ES: Equal Share + SJF
Keep filling the available resources

SRPT: Shortest Remaining Processing Time
Impractical switching between CPU&GPU

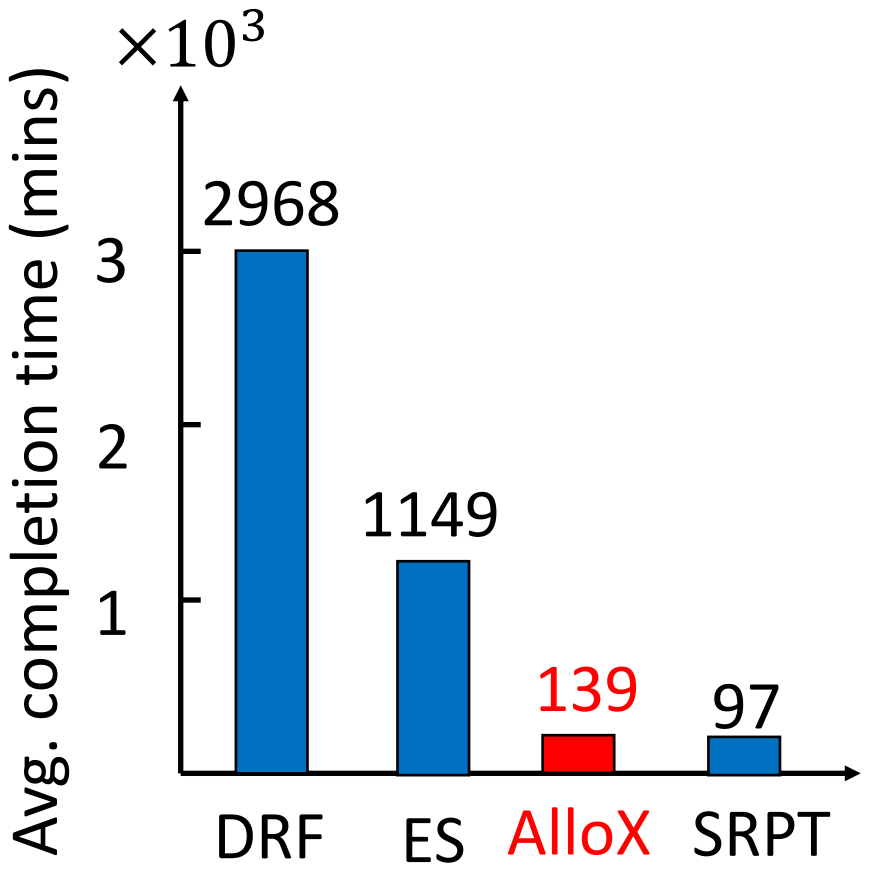

AlloX reduces up to 95% avg. completion time

TensorFlow CNN benchmarks

AlloX: Compute Allocation in Hybrid Clusters

Tan N. Le Xiao Sun Mosharaf Chowdhury Zhenhua Liu

Stony Brook University

SUNY Korea The State University of New York

UNIVERSITY OF MICHIGAN

tnle@cs.stonybrook.edu